

Improving Energy Efficiency in Upstream EPON Channels by Packet Coalescing

Miguel Rodríguez-Pérez, Sergio Herrería-Alonso,
Manuel Fernández-Veiga, *Member, IEEE*, and Cándido López-García

Abstract—In this paper, we research the feasibility of adapting the packet coalescing algorithm, used successfully in IEEE 802.3az Ethernet cards, to upstream EPON channels. Our simulation experiments show that, using this algorithm, great power savings are feasible without requiring any changes to the deployed access network infrastructure nor to protocols.

Index Terms—Energy efficiency, EPON, packet coalescing

I. INTRODUCTION

THE TOTAL amount of energy needed to power networking infrastructure has been rising quickly as more and faster devices get connected to the Internet. In the last few years, these growing power demands have entered in conflict with the public awareness about environmental sustainability and the higher operating costs associated with the networking machinery.

In this paper, we focus on reducing the power consumption of Ethernet passive optical networks (EPONs), that are among the most popular access network systems nowadays. Although EPONs are power efficient when operating at 1 Gb/s [1], their energy needs grow significantly when the line rate is increased to 10 Gb/s [2], so an energy saving mechanism should be employed to reduce power consumption.

Basically, an EPON consists of one optical line terminal (OLT), located at the provider central office, connected via optical fibers and optical passive splitters to multiple optical network units (ONUs) located at the users' premises. In downstream transmission, data are broadcasted from the OLT to all the connected ONUs, so each ONU must filter out those packets not directed to itself. In the upstream direction, a common upstream channel is shared among all the ONUs using time division multiple access (TDMA). The OLT allocates the appropriate share of upstream bandwidth to each ONU with the help of a dynamic bandwidth allocation (DBA) algorithm that takes into account their different needs. Additionally, the multi-point control protocol (MPCP) is used as the MAC algorithm to emulate a dedicated point-to-point channel from each ONU to the OLT.

During the last years, several mechanisms to reduce the power consumption of EPONs have appeared. Unfortunately, most of them require the modification of the MPCP protocol [3], [4], [5], [6] or rely on a fixed bandwidth allocation

(FBA) that may result in bandwidth under or over-allocation, since fixed time slots are assigned to each ONU without considering their diverse bandwidth demands [7].

In this paper, we present a promising mechanism to save energy in the ONUs allowing them to enter a low power mode when they do not need to send traffic. We will consider that, when an ONU enters this sleep mode, its upstream transmission circuitry can be powered down and thus, its energy demands are minimal, though, obviously, no transmission can be carried out and any upstream traffic must be buffered.¹ Our proposal does not need any modification to the MPCP protocol or the DBA algorithm. We based our work on the *packet coalescing* algorithm successfully applied to reduce power consumption in Ethernet interfaces [8], [9]. Thus, instead of waking up the ONU in the presence of upstream traffic as in [10], [6], we propose to delay the exit from the sleep mode until the upstream queue reaches a certain threshold. Bringing the technique from the Ethernet domain to EPON channels is not straightforward. The MPCP protocol constraints the ONU restricting the transmission opportunities and demands anticipation on its part. In this paper we propose a new state machine for sleep transitions at the ONU that permits to use the packet coalescing algorithm even on MPCP regulated channels. We conduct several simulation experiments to study the effectiveness of this mechanism in terms of energy consumption and queuing delay. Simulation results show that our proposal is able to provide great energy savings at the expense of slightly increasing the queueing delay.

II. DESCRIPTION

A. The DBA cycle

In essence, under MPCP the upstream channel is divided in periods that ONUs employ to report their upstream queue lengths and to transmit their traffic to the OLT. These periods correspond with DBA cycles. Each period the OLT sends a *gate message* to every ONU containing two time intervals for the next DBA cycle: One time interval to be used to transmit data from the ONU to the OLT and the other for sending a *traffic report*. In these reports, the ONUs indicate the amount of data stored in their upstream queues. Then, the DBA algorithm allocates upstream capacity to each ONU based on the reports received in the previous cycle. Thus, an ONU that had no traffic to send in DBA cycle i , and hence

The authors are with the Telematics Engineering Dept., Univ. of Vigo, 36310 Vigo, Spain. Tel.: +34 986 813458; fax: +34 986 812116; email: miguel@det.uvigo.es (M. Rodríguez-Pérez). This work was supported by the "Ministerio de Ciencia e Innovación" through the project TEC2009-12135 of the "Plan Nacional de I+D+I" (partly financed with FEDER funds).

¹As we are focused in upstream channel savings, we consider ONUs with separate upstream/downstream circuitry, that can be powered down independently.




Figure 1. State machine diagram of the upstream packet coalescing algorithm.




Figure 2. Time line of OFF to ON transition and back.

reported empty queues to the OLT, will not receive upstream capacity in cycle $i + 1$.

When an ONU has no upstream data to transmit, it can spare to send its report packet if the previous one was sent less than 50 ms earlier. However, if the OLT does not receive a report from an ONU for more than 50 ms, the ONU is considered disconnected by the OLT. This imposes a clear upper limit on the time an ONU can remain continuously sleeping.

B. Upstream Channel Power Saving

With the help of the diagram in Fig. 1 we will describe the state machine of our sleep capable ONU. Under our proposal, an ONU remains in the OFF state until the upstream queue length q surpasses a certain threshold q_w or the time since the previous report was sent reaches 50 ms. However, the ONU is not yet ready to transmit immediately after the OFF state is left, and thus a new transient WAIT state is entered. While in WAIT, the transmission circuitry remains off until the proper time to switch it on arrives, as explained below.

Before being allowed to transmit its queued data, the ONU must first send a report to the OLT to be allocated a transmission slot in the future DBA cycle. However, the ONU cannot send its report immediately, but must wait until the time indicated in the previous gate message received from the OLT. As the time needed to power up the transmitting circuitry (δ_{ON}) is not null, the ONU remains in the WAIT state until δ_{ON} seconds before the start of the next DBA cycle, in which it will be able to transmit its report. Then, it switches to the TRANS state. Note that it is upon entering this state that the transmission circuitry is powered on. The timing of this procedure can also be seen in Fig. 2.

The ONU stays in TRANS until it is finally allowed to start transmitting data. For this it must wait for a whole DBA cycle. At the start of the cycle it will report its queue length to the OLT.² Later in the cycle, the OLT will send a new gate

²The exact time will have been indicated in the gate message received in the previous DBA cycle.

Table I
SIMULATION PARAMETERS

Frame size	1 500 bytes
Arrivals distribution	Pareto ($\alpha = 2.5$)
Upstream queue capacity	∞
Available upstream bandwidth	200 Mb/s
Nominal upstream bandwidth	10 Gb/s
DBA cycle length	1.5 ms
Active to sleep power ratio	10 : 1
δ_{ON}	2 ms
q_w	1, 10 and 100

message containing both the transmission time of the next report and the slot for data transmission, both to occur in the next DBA cycle. So, after finishing this transitioning DBA cycle the ONU will proceed to the *normal* ON state. Note that the TRANS state is different from the ON one in that, despite the fact that the ONU is completely powered on in both of them, it only transmits data in ON, while in TRANS it does no useful work. In fact, only administrative traffic (report messages) gets transmitted.

The transition from ON to OFF is somewhat simpler. Once the ONU empties its upstream transmission queue, it remains ON until it sends a final report in the next DBA cycle to indicate that the queue is empty.³ Only when the report is finally sent, and the upstream queue remains below the q_w threshold,⁴ the ONU switches off the transmission circuitry and enters the OFF state. We consider the time needed to power down the circuitry to be negligible.

III. RESULTS

In this section we provide results for our proposed use of *packet coalescing* in the upstream channel of GEAPON networks. We made use of an in-house simulator, available for download at [11]. For simplicity the simulator assumes a number of conditions. Mainly that the duration of the DBA cycle is fixed and that the upstream bandwidth available to each ONU is capped to a constant value. Upstream traffic follows a Pareto distribution ($\alpha = 2.5$). We used 200 Mb/s for the available upstream bandwidth per ONU, as a kind of worst case scenario with fifty ONUs demanding maximum bandwidth in a 10 Gb/s link. The rest of the experiment parameters are detailed in Table I. The transition time has been set to the same value as in [6], [4], [12]. The power ratio between the active and sleep status $P_{ON}/P_{OFF} = 10$ has been chosen based on typical values in the literature [3], [4], [6].⁵

In the first experiment we have measured the time spent in each state for different values of the maximum queuing

³Another possibility is for the ONU to shut down immediately without sending this last report. Although this would increment power savings, we opted for a more conservative approach for the sake of interoperability.

⁴If the design opted for a more stringent approach, that is, demanding the queue to be completely empty when entering OFF state, the ONU would be prevented to save energy in some scenarios. The time between the end of the transmission and the final report is not negligible, and a few packets may arrive in between. If the ONU did not enter OFF in spite of those packets, it would transmit them in the next DBA cycle, defeating the packet coalescing strategy, and thus reducing energy efficiency.

⁵We consider the power consumption of the WAIT state to be equal to that of the OFF state, as the transmission circuitry is powered off in both states. Similarly, the power needs of the TRANS and ON states are also the same.

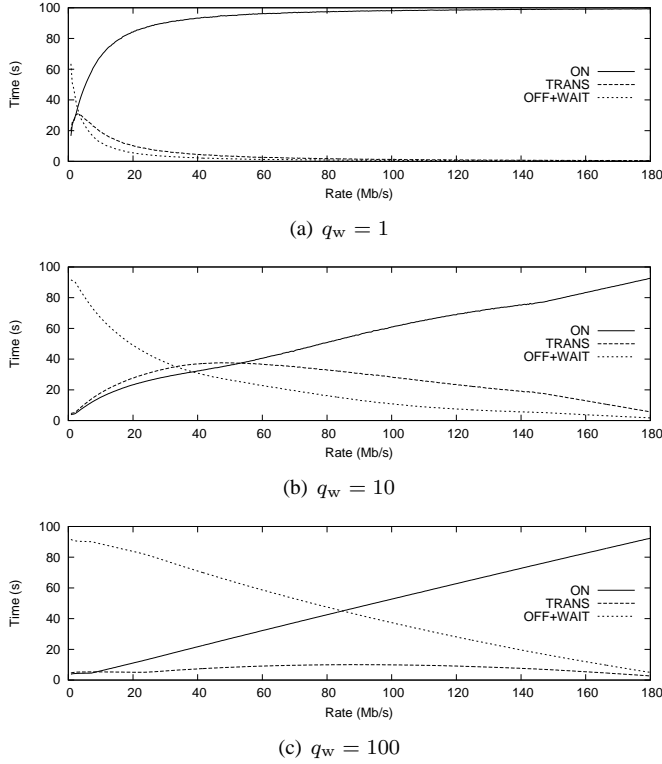


Figure 3. Total time on each state for a 10 Gb/s EPON link with 200 Mb/s average upstream capacity.

threshold. Each simulation was run for 100 seconds and repeated with different random seeds.⁶

Figure 3(a) shows the results obtained for $q_w = 1$. As expected, the time spent ON grows with the offered load while, conversely, the sleeping time decreases. Note, however, how the time wasted in the TRANS state represents a significant share specially for low loads, where a power saving algorithm should be able to save more power. This is because, every time a new packet arrives while sleeping, TRANS is entered wasting a fixed amount of time until the ON state is finally reached and the backlogged traffic (a few packets, usually) is transmitted. For moderate loads the ONU is likely ON because the probability for new packet arrivals while waiting to send the final report and enter OFF is very high.

If the q_w threshold is set to 10 packets we obtain the results of Fig. 3(b). Now the time spent in ON and OFF+WAIT becomes somewhat more linear. The time in TRANS still represents a significant share, but the time ON clearly diminishes. Now there is less of a chance for 10 packets to arrive while waiting to switch off the circuitry after a transmission has emptied the queue.

Figure 3(c) shows the results when q_w is raised to 100 packets. This kind of extreme case shows, as expected, the best behavior. Now the time in TRANS is very low and the ON and OFF curves are almost linear. The TRANS time decreases because, for the same amount of traffic, the number of needed transitions to ON is less than before, as the ONU

⁶95 % confidence intervals were negligible and are not represented so as not to clutter the figures.




Figure 4. Power consumption in a 10 Gb/s EPON link with 200 Mb/s average upstream capacity for different q_w thresholds.




Figure 5. Average packet delay in a 10 Gb/s EPON link with 200 Mb/s average upstream capacity for different q_w thresholds.

waits to accumulate more upstream traffic before exiting OFF. Moreover, the time spent ON is more wisely employed, as more DBA cycles are usually needed to transmit all the accumulated traffic. This also makes the time used to wait for the final report less important when compared to the total time in the ON state. That is, now the majority of the time in the ON state is employed to actually transmit traffic.

The time spent in each status has a direct impact on the total amount of power consumed. Figure 4 shows the total amount of power used for the different values of q_w . Obviously, a non power-aware ONU would always consume 100 % of power, while an ideal power saving mechanism would show a linear increase in consumption. In the figure we observe how the packet coalescing algorithm gets very good results for both $q_w = 10$ and $q_w = 100$, with the latter being really close to the ideal power saving mechanism. With $q_w = 1$ however, the algorithm gets badly punished by the frequent state transitions that waste power while not transmitting traffic.

Our latest experiment measures the impact of the different possible configurations of the algorithm on average packet delay. The results are plotted in Fig. 5. As expected, the added delay increases for the greatest values of q_w . In any case, the added delay stays always below 35 ms, a delay not unheard of in other last mile technologies, like xDSL. Moreover, for low loads, where the delay is maximum, it does not depend as much on the value of q_w as on the maximum time in the OFF state. This is why both $q_w = 10$ and $q_w = 100$ converge for low loads. If necessary, the maximum time while sleeping can be reduced, albeit with effects in power savings.

IV. CONCLUSIONS

In this paper we have successfully adapted the *packet coalescing* algorithm originated in Ethernet cards to the upstream data channel of EPON networks. The adaptation had to take into account the restrictions imposed by the MPCP MAC protocol that regulates the access to the shared upstream channel among the different ONUs. Simulation results show

that our proposal is able to achieve great power savings, very close in fact to the optimum, with only a bounded moderate increase in network delay.

Finally, we are also working on coordinating the power savings in the upstream channel with those that can be achieved if the downstream one could also be put to sleep. This will probably require modifications to MPCP or the cooperation of the OLT as it will have to restrict *gate messages* to the periods where the downstream interface at the ONU is ready. A mechanism to calculate (and coordinate) optimum sleeping periods, like those present in [13] or [14], will probably be useful.

REFERENCES

- [1] C. Lange and A. Gladisch, "On the energy consumption of FTTH access networks," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2009.
- [2] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical IP networks," *J. Lightw. Technol.*, vol. 27, no. 13, pp. 2391–2403, Jul. 2009.
- [3] D. Ren, H. Li, and Y. Ji, "Power saving mechanism and performance analysis for 10 gigabit-class passive optical network systems," in *Network Infrastructure and Digital Content, 2010 2nd IEEE International Conference on*, Sep. 2010, pp. 920–924.
- [4] J. Mandin, "EPON powersaving via sleep mode," IEEE P802.3av 10GEPON Task Force Meeting, Sep. 2008.
- [5] Y. Yan, S.-W. Wong, L. Valcarenghi, S.-H. Yen, D. R. Campelo, S. Yamashita, L. Kazovsky, and L. Dittmann, "Energy management mechanism for Ethernet passive optical networks (EPONs)," in *IEEE Int. Conf. on Comm. (ICC)*, May 2010.
- [6] R. Kubo, J. Kani, H. Ujikawa, T. Sakamoto, Y. Fujimoto, N. Yoshimoto, and H. Hadama, "Study and demonstration of sleep and adaptive link rate control mechanisms for energy efficient 10G-EPON," *J. of Opt. Comm. and Netw.*, vol. 2, no. 9, pp. 716–729, Aug. 2010.
- [7] S. S. W. Lee and A. Chen, "Design and analysis of a novel energy efficient Ethernet passive optical network," in *9th Int. Conf. on Networks*, Apr. 2010, pp. 6–9.
- [8] K. Christensen, P. B. N. r. Reviriego Vasallo, M. Bennett, M. Mostowfi, and J. A. Maestro, "IEEE 802.3az: the road to energy efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, no. 11, pp. 50–56, 2010.
- [9] S. Herrería Alonso, M. Rodríguez Pérez, M. Fernández Veiga, and C. López García, "A power saving model for burst transmission in energy-efficient Ethernet," *IEEE Commun. Lett.*, vol. 15, no. 5, pp. 584–586, May 2011.
- [10] J. Zhang and N. Ansari, "Toward energy-efficient 1G-EPON and 10G-EPON with sleep-aware MAC control and scheduling," *IEEE Commun. Mag.*, vol. 49, no. 2, pp. s33–s38, Feb. 2011.
- [11] M. Rodríguez Pérez, "A GEPON simulator," <http://labredes.det.uvigo.es/~miguel/gepon>, Jun. 2011.
- [12] S.-W. Wong, L. Valcarenghi, S.-H. Yen, D. R. Campelo, S. Yamashita, and L. Kazovsky, "Sleep mode for energy saving PONs: Advantages and draw backs," in *IEEE GLOBECOM Workshops*, Honolulu, HI, USA, Dec. 2009, pp. 1–6.
- [13] M. Gupta and S. Singh, "Using low-power modes for energy conservation in Ether net LANs," in *Proceedings of the IEEE INFOCOM*, 2007, pp. 2451–2455.
- [14] M. Rodríguez Pérez, S. Herrería Alonso, M. Fernández Veiga, and C. López García, "Improved opportunistic sleeping algorithms for LAN switches," in *Proceedings of the IEEE Globecom*, Honolulu, HI, USA, Dec. 2009.